arXiv:1609.05692v2 [q-bio.QM] 5 Nov 2016

Supplementary information

Improving landscape inference by integrating heterogeneous data in the inverse Ising problem

Pierre Barrat-Charlaix, Matteo Figliuzzi, and Martin Weigt

Details of the biological data

The following datasets are used in the analysis of biological data:

- **Multiple sequence alignements** We used multiple sequence alignments (MSA) of sequences belonging to the Pfam Beta-lactamase2 family PF13354 for the inference of TEM-1 mutational landscape (resp. PDZ domain family PF00595 for PSD-95 mutational landscape) [30]. Raw data have been processed as described in [18]: we have used HMMer [29] to search against the Uniprot protein sequence database (version of March 2015); from the resulting alignments we have removed all sequences with more than 5 gapped positions; in order to take into account phylogenetic correlations and sampling biases we have associated to each sequence $\boldsymbol{s^m}$ the following statistical weight:

$$w_m = \left(1 + \sum_{m \neq m'} \Theta\Big(d_{m,m'} - \vartheta N\Big)\right)^{-1} , \tag{1}$$

  with $d_{mm'}$ being the Manhattan distance (number of mismatches) between sequences $\boldsymbol{s^m}$ and $\boldsymbol{s^{m'}}$ and Θ being the Heaviside step function whose value is zero for negative argument and one for positive argument. The reweighting threshold is set to $\vartheta = 0.8$ as usually done in DCA inference [32]. The resulting MSA are $N = 197$ (resp. $N = 81$) sites long, and contain $M = 2462$ (resp. $M = 3287$) homologous sequences (average sequence identity $\sim 20\%$ with the wild-type sequences).

- **Computational predictions of thermodynamic stabilities** We used $\Delta\Delta G$-predictions for all single amino acid mutations from the wild-type sequences as estimated by PoPMuSic web-based tool [31] taking as reference structure the PDB: 1M40 for TEM-1 (resp PDB: 1BE9 for PSD-95).

- **Experimental fitness measurements**
  - **TEM-1.** Experimental measurements from [23] provide estimations of amoxicillin Minumum Inibitory Concentration MIC for $n = 742$ distinct single amino-acid mutations falling inside the part of the sequence covered by the corresponding Pfam model. Following [18] we have defined the experimental fitness as the logarithmic average over all MIC measurements of a specific mutant sequence (whenever multiple measurements for the same sequence were available).

– **PSD-95.** In [22] the effect of each single point mutations has been quantified as the log frequency of observing each amino acid β at each position i in the selected (sel) versus the unselected (unsel) population, relative to amino acid s_i present in the wild type:

$$\phi(s_i \to \beta) = \log\left[\frac{f_i^{sel}(\beta)}{f_i^{sel}(s_i)}\right] - \log\left[\frac{f_i^{unsel}(\beta)}{f_i^{unsel}(s_i)}\right] \tag{2}$$

Potts Hamiltonian

The relevant description in the statistical modeling of protein sequences is a 21-state Potts Model, since each variable $s_i, i = 1, ..., N$, can now assume 21 states (20 amino acids, one alignment gap). The following Potts Hamiltonian provides the energy of any amino-acid sequence $\boldsymbol{s} = (s_1, ..., s_N)$,

$$\mathcal{H}^{\text{Potts}}(\boldsymbol{s}) = -\sum_{i<j}^{N} J_{ij}(s_i, s_j) - \sum_{i=1}^{N} h_i(s_i) \ , \tag{3}$$

where fields h_i and couplings J_{ij} are now respectively 21-dimensional vectors and 21×21-dimensional matrices. We fix the gauge freedom (cf. [1]) using, for all i, j, the lattice-gas conditions $h_i(s_i^{ref}) = 0$ and $J_{ij}(s_i^{ref}, s_j) = J_{ij}(s_i, s_j^{ref}) = 0$. Here $\boldsymbol{s}^{ref} = (s_1^{ref}, ..., s_N^{ref})$ is the reference wildtype sequence, for which mutations shall be predicted.

Fitness-energy mapping

Since model energies $\mathcal{H}(\boldsymbol{\sigma^a})$ can be in any non-linear relation with measured fitnesses ϕ^a we replace Eq. (9) with:

$$\begin{aligned} p_{i,\alpha}(\mathbf{J}, \mathbf{h}) &= \pi_{i,\alpha}^{\text{emp}} + \frac{\lambda}{1-\lambda}\frac{1}{M}\sum_{a=1}^{P} \delta_{\sigma_i^a,\alpha}\left[\phi^a - \mu\big(\mathcal{H}(\boldsymbol{\sigma}^a)\big)\right] \\ p_{ij,\alpha\beta}(\mathbf{J}, \mathbf{h}) &= \pi_{ij,\alpha\beta}^{\text{emp}} + \frac{\lambda}{1-\lambda}\frac{1}{M}\sum_{a=1}^{P} \delta_{\sigma_i^a,\alpha}\delta_{\sigma_j^a,\beta}\left[\phi^a - \mu\big(\mathcal{H}(\boldsymbol{\sigma}^a)\big)\right] \end{aligned} \tag{4}$$

where $p_{i,\alpha}(\mathbf{J}, \mathbf{h})$ (resp. $\pi_{i,\alpha}^{\text{emp}}$) are the single marginals in the model (resp. in the empirical sample) specyfing the frequency of observing amino-acid α at position i, $p_{ij,\alpha\beta}(\mathbf{J}, \mathbf{h})$ (resp. $\pi_{ij,\alpha\beta}^{\text{emp}}$) are the pairwise marginals, and $\mu(x)$ is the robust monotonous non-linear mapping from energy to fitness introduced in [18]. The mapping is obtained by first sorting the model energies and experimental fitnesses and then associating the n_{th} smallest energy $\mathcal{H}(n_{th})$ with the n_{th} highest experimental fitness value $\phi(n_{th})$,

$$\mu\big(\mathcal{H}(n_{th})\big) = \phi(n_{th}) \ . \tag{5}$$

Note that the second terms in Eq. (4) vanish when model energies are exactly in the inverse order of experimental fitnesses. We subsequently used the mapping to compute linear correlations between the mapped energies $\mu\big(\mathcal{H}(\boldsymbol{\sigma}^a)\big)$ and the experimental fitnesses ϕ^a, resulting in non-linear rank correlations between experimental fitnesses and model energies.

arXiv:1609.05692v2 [q-bio.QM] 5 Nov 2016

Improving landscape inference by integrating heterogeneous data in the inverse Ising problem

Pierre Barrat-Charlaix[a,+], **Matteo Figliuzzi**[a,b,+], **and Martin Weigt**[a,*]

[a]Sorbonne Universités, UPMC Univ Paris 06, CNRS, Biologie computationnelle et quantitative - Institut de Biologie Paris Seine, 75005 Paris, France
[b]Sorbonne Universités, UPMC Univ Paris 06, Institut de Calcul et de la Simulation, 75005 Paris, France
[+]these authors contributed equally to this work
[*]To whom correspondence should be addressed. Email: martin.weigt@upmc.fr

ABSTRACT

The inverse Ising problem and its generalizations to Potts and continuous spin models have recently attracted much attention thanks to their successful applications in the statistical modeling of biological data. In the standard setting, the parameters of an Ising model (couplings and fields) are inferred using a sample of equilibrium configurations drawn from the Boltzmann distribution. However, in the context of biological applications, quantitative information for a limited number of microscopic spins configurations has recently become available. In this paper, we extend the usual setting of the inverse Ising model by developing an integrative approach combining the equilibrium sample with (possibly noisy) measurements of the energy performed for a number of arbitrary configurations. Using simulated data, we show that our integrative approach outperforms standard inference based only on the equilibrium sample or the energy measurements, including error correction of noisy energy measurements. As a biological proof-of-concept application, we show that mutational fitness landscapes in proteins can be better described when combining evolutionary sequence data with complementary structural information about mutant sequences.

Introduction

High-dimensional data characterizing the collective behavior of complex systems are increasingly available across disciplines. A global statistical description is needed to unveil the organizing principles ruling such systems and to extract information from raw data. Statistical physics provides a powerful framework to do so. A paradigmatic example is represented by the Ising model and its generalizations to Potts and continuous spin variables, which have recently become popular for extracting information from large-scale biological datasets. Successful examples are as different as multiple-sequence alignments of evolutionary related proteins[1–3], gene-expression profiles[4], spiking patterns of neural networks[5,6], or the collective behavior of bird flocks[7]. This widespread use is motivated by the observation that the least constrained (i.e. maximum-entropy[8]) statistical model reproducing empirical single-variable and pairwise frequencies observed in a list of equilibrium configurations is given by a Boltzmann distribution:

$$P(s) = \frac{1}{\mathscr{Z}} \exp\{-\mathscr{H}(s)\}, \quad \mathscr{H} = -\sum_{i<j}^{N} J_{ij} s_i s_j - \sum_{i=1}^{N} h_i s_i \,, \tag{1}$$

with $s = (s_1, \ldots, s_N)$ being a configuration of N binary variables or 'spins'. Inferring the couplings $\mathbf{J} = \{J_{ij}\}_{1\le i<j\le N}$ and fields $\mathbf{h} = \{h_i\}_{1\le i\le N}$ in the Hamiltonian $\mathscr{H}$ from data, known as the *inverse Ising problem*, is computationally hard for large systems ($N \gg 1$). It involves the calculation of the partition function $\mathscr{Z} = \sum_s e^{-\mathscr{H}(s)}$ as a sum over an exponential number of configurations. The need to develop efficient approximate approaches has recently triggered important work within the statistical-physics community, cf. e.g.[9–17].

Despite the broad interest in inverse problems, the methodological setting has remained rather limited: all of this literature, including the biological cases mentioned in the beginning, seeks to estimate model parameters starting from a set of configurations s, which are considered to be at equilibrium and independently drawn from $P(s)$. Real data, however, may be quite different. In biological systems, "microscopic spins configurations" (e.g. amino-acid sequences) are increasingly accessible to experimental techniques, and *quantitative information* for a limited number of *particular* configurations (e.g. three-dimensional structures, measured activities or thermodynamic stabilities for selected proteins) is frequently available. It seems reasonable to actually integrate such information into the inverse Ising problem instead of ignoring it. In this work, we use two different types of data (cf. Fig. 1):

- As in the standard inverse Ising problem, part of the data comes as a sample of equilibrium configurations assumed to be drawn from the Boltzmann distribution to be inferred.
- The second data source is a collection of arbitrary configurations together with *noisy* measurements of their energy.

These data sets are limited in size and accuracy. Therefore an optimized integration of both data types is expected to improve the overall performance as compared to the individual use of one single data set.

The inspiration to develop this new integrative framework for the inverse Ising problem is taken from *protein fitness landscapes* in biology, which provide a quantitative mapping from any amino-acid sequence $s = (s_1, ..., s_N)$ to a fitness $\phi(s)$ measuring the ability of the corresponding protein to perform its biological function. Fitness landscapes are of outstanding importance in evolutionary and medical biology, but it appears impossible to deduce a protein's fitness from its sequence only. Experimental or computational approaches exploiting other data are urgently needed.

Information about fitness landscapes can be found in the amino-acid statistics observed in natural protein sequences, which are related to the protein of interest. In fact they represent diverse but functional configurations sampled by evolution. It has been recently proposed that their statistical variability can be captured by Potts models (generalization of the Ising model to 21-state amino-acid variables). Indeed, statistical models inferred from large collections of natural sequences have recently led to good predictions of experimentally measured effects[18–21]: in a number of systems, the fitness cost $\Delta\phi(s) \equiv \phi(s) - \phi(s^{\rm ref})$ of mutating any amino acid in a reference protein $s^{\rm ref}$ strongly correlates with the corresponding energy changes in the inferred statistical model,

$$\Delta\phi(s) \sim \quad log\left[\frac{P(s)}{P(s^{\rm ref})}\right] = -\big(\mathscr{H}(s) - \mathscr{H}(s^{\rm ref})\big), \tag{2}$$

suggesting that the Hamiltonian of the inferred models is strictly related to the underlying mutational landscapes.

While evolutionary diverged sequences can be regarded as a global sample of the fitness landscape, further information can be obtained from direct measurements on particular *'microstates'* of the system, *i.e.* individual protein sequences. Recent advances in experimental technology allow for conducting large-scale *mutagenesis* studies: in a typical experiment, a reference protein of interest is chosen, and a large number ($10^3 - 10^5$) of mutant proteins (having sequences differing by one or few amino acids from the reference) are synthesized and then characterized in terms of fitness. This provides a systematic *local* measurement of the fitness landscape[22–24]. Regression analysis may be used to globally model mutational landscapes[25]. A second-order parameterization of ϕ arises naturally in this context, when considering an expansion of effects in terms of independent additive effects and pairwise *'epistatic'* interactions between sites[26],

$$\phi(s) = \sum_{i=1}^{N} \varphi_i(s_i) + \sum_{1\le i\le j\le N} \varphi_{i,j}(s_i, s_j)\ . \tag{3}$$

However, the number of accessible mutant sequences remains small compared to the number of terms in this sum, and mutagenesis data alone are not sufficient to faithfully model fitness landscapes[27].

In situations where no single dataset is sufficient for accurate inference, integrative methods accounting for complementary data sources will improve the accuracy of computational predictions. In this paper, (i) we define a generalized inference framework based on the availability of an equilibrium sample *and* of complementary quantitative information; (ii) we propose a Bayesian integrative approach to improve over the limited accuracy obtainable using standard inverse problems; (iii) we demonstrate the practical applicability of our method in the context of predicting mutational effects in proteins, a problem of outstanding bio-medical importance for questions related to genetic disease and antibiotic drug resistance.

Results

An integrated modeling

The inference setting. Inspired by this discussion, we consider two different datasets originating from a *true* model $\mathscr{H}^0$. The first one, $\mathbf{D}_{\rm eq} = \{s^1, ..., s^M\}$, is a collection of M equilibrium configurations independently drawn from the Boltzmann distribution $P^0(s)$. For simplicity, we consider binary variables $s_i \in \{0, 1\}$, corresponding to a "lattice-gas" representation of the Ising model in Eq. (1). This implies that energies are measured with respect to the reference configuration $s^{\rm ref} = (0, ..., 0)$. The standard approach to the inverse Ising problem uses only this type of data to infer parameters of $\mathscr{H}$: couplings and fields in Eq. (1) are fitted so that the inferred model reproduces the empirical single and pairwise frequencies,

$$\pi_i^{\rm emp} = \frac{1}{M}\sum_{\mu=1}^{M} s_i^\mu, \qquad \pi_{ij}^{\rm emp} = \frac{1}{M}\sum_{\mu=1}^{M} s_i^\mu s_j^\mu\ . \tag{4}$$

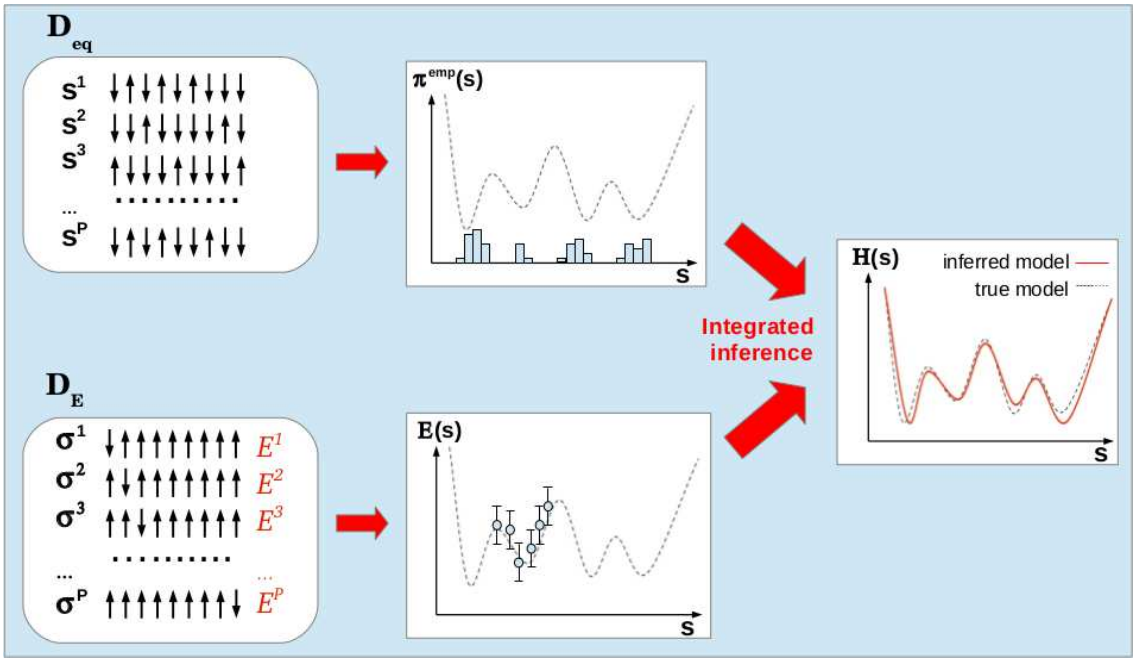

Figure 1. Schematic representation of the inference framework: a sample of equilibrium configurations (dataset $\mathbf{D}_{\rm eq}$) and noisy energy measurements for another set of configurations (dataset $\mathbf{D}_{\rm E}$) are integrated within a Bayesian approach to infer the model $\mathscr{H}$. The dashed lines represent the underlying true landscape, which has to be inferred, the red line the inferred landscape.

The second dataset provides a complementary source of information, which shall be modeled as noisy measurements of the *energies* of a set of P *arbitrary* (i.e. not necessarily equilibrium) configurations σ^a. These data are collected in the dataset $\mathbf{D}_E = \{(E^a, \sigma^a)_{a=1,\dots,P}\}$, with

$$E^a = \mathscr{H}^0(\sigma^a) + \xi^a \quad a = 1, \dots, P \,. \tag{5}$$

The noise ξ^a models measurement errors or uncertainties in mapping measured quantities to energies of the Ising model. For simplicity, we consider ξ^a to be white Gaussian noise with zero mean and variance Δ^2: $\langle \xi^a \xi^b \rangle = \delta_{a,b}\Delta^2$.

As schematically represented in Fig. 1, datasets $\mathbf{D}_{\rm eq}$ and $\mathbf{D}_E$ constitute different sources of information about the energy landscape defined by Hamiltonian $\mathscr{H}^0$. Observables in Eq. (4) are empirical averages computed from equilibrium configurations in $\mathbf{D}_{\rm eq}$, providing *global* information about the energy landscape. On the contrary, configurations in $\mathbf{D}_E$ are arbitrarily given, and a (noisy) measurement of their energies provides *local* information on particular points in the landscape.

A maximum-likelihood approach. To infer the integrated model, we consider a *joint description of the probabilities of the two data types* for given parameters $\mathbf{J}$ and $\mathbf{h}$ of tentative Hamiltonian $\mathscr{H}$. The probability of observing the sampled configurations in $\mathbf{D}_{\rm eq}$ equals the product of the Boltzmann probability of each configuration,

$$P(\mathbf{D}_{\rm eq}|\mathbf{J},\mathbf{h}) = \exp\left\{ -\sum_{\mu=1}^{P} \mathscr{H}(\mathbf{s}^\mu) - M \log \mathscr{Z}(\mathbf{J},\mathbf{h}) \right\} . \tag{6}$$

To derive an analogous expression for the second dataset, we integrate over the Gaussian distribution of the noise $\xi^a = E^a - \mathscr{H}(\sigma^a)$ obtaining a Gaussian probability of the energies (remember configurations in $\mathbf{D}_E$ are arbitrarily given):

$$\begin{aligned} P(\mathbf{D}_E|\mathbf{J},\mathbf{h}) &= \prod_{a=1}^{P} \int d\xi_a P(\xi_a)\, \delta(E^a - \mathscr{H}(\sigma^a) - \xi^a) \\ &= \frac{1}{(2\pi\Delta^2)^{\frac{P}{2}}} \exp\left\{ -\sum_{a=1}^{P} \frac{[E^a - \mathscr{H}(\sigma^a)]^2}{2\Delta^2} \right\} \end{aligned} \tag{7}$$

The combination of these expressions provides the joint log-likelihood for the model parameters given the data:

$$\mathscr{L}(\mathbf{J},\mathbf{h}|\mathbf{D}_{\rm eq},\mathbf{D}_E) = \log P(\mathbf{D}_{\rm eq}|\mathbf{J},\mathbf{h}) + \log P(\mathbf{D}_E|\mathbf{J},\mathbf{h}) \tag{8}$$

Maximizing the above likelihood with respect to parameters $\{h_i\}_{1\leq i\leq N}$ and $\{J_{ij}\}_{1\leq i<j\leq N}$ leads to the following self-consistency equations:

$$\begin{aligned} p_i(\mathbf{J},\mathbf{h}) &= \pi_i^{\text{emp}} + \frac{\lambda}{1-\lambda}\frac{1}{M}\sum_{a=1}^{P}\sigma_i^a\left[E^a - \mathscr{H}(\sigma^a)\right] \\ p_{ij}(\mathbf{J},\mathbf{h}) &= \pi_{ij}^{\text{emp}} + \frac{\lambda}{1-\lambda}\frac{1}{M}\sum_{a=1}^{P}\sigma_i^a\sigma_j^a\left[E^a - \mathscr{H}(\sigma^a)\right] \end{aligned} \tag{9}$$

with $p_i(\mathbf{J},\mathbf{h}) = \langle\sigma_i\rangle_{\mathscr{H}}$, $p_{ij}(\mathbf{J},\mathbf{h}) = \langle\sigma_i\sigma_j\rangle_{\mathscr{H}}$ being single and pairwise averages in the model Eq. (1). We have introduced the parameter $\lambda = \frac{1}{1+\Delta^2}$: in practical applications, the error Δ may not be known, and the parameter $0 \leq \lambda < 1$ allows to weigh data sources differently. For $\lambda = 0$ (*i.e.* large noise), the standard inverse Ising problem is recovered: optimal parameters are such that the model exactly reproduces magnetizations and correlations of the sample. For $\lambda > 0$, the second dataset containing quantitative data is taken into account: whenever energies computed from the Hamiltonian $\mathscr{H}$ do not match the measured ones, the model statistics deviates from the sample statistics. Both log-likelihood terms in (8) are concave, and thus their sum: Eq. (9) has a unique solution.

Noiseless measurements. The case of noiseless energy measurements in Eq. (5) (i.e. $\lambda \to 1$) has to be treated separately. First, energies have to be perfectly fitted by the model, by solving the following linear problem:

$$\mathbf{X}\vec{\mathscr{H}} = \vec{E}\ , \tag{10}$$

where we have introduced a $N(N+1)/2$ dimensional vectorial representation $\vec{\mathscr{H}} = (h_1,...,h_N,J_{12},...,J_{N,N-1})^T$ of the model parameters, and $\vec{E} = (E_1,...,E_P)^T$ contains the exactly measured energies. The matrix

$$\mathbf{X} = \begin{bmatrix} \sigma_1^1 & \cdots & \sigma_N^1 & (\sigma_1^1\sigma_2^1) & \cdots & (\sigma_{N-1}^1\sigma_N^1) \\ \vdots & & \vdots & \vdots & & \vdots \\ \sigma_1^P & \cdots & \sigma_N^P & (\sigma_1^P\sigma_2^P) & \cdots & (\sigma_{N-1}^P\sigma_N^P) \end{bmatrix} \tag{11}$$

specifies which parameters contribute to the energies of configurations in the second dataset. If $K = N(N+1)/2 - rank(X) > 0$, the parameters cannot be uniquely determined from the measurements: The sample $\mathbf{D}_{\text{eq}}$ can be used to remove the resulting degeneracy. To do so, we parametrize the set of solutions of Eq. (10) as follows:

$$\vec{\mathscr{H}} = \vec{\mathscr{H}}_{nh} + \sum_{k=1}^{K}\alpha_k\vec{\mathscr{O}}_k \tag{12}$$

where $\vec{\mathscr{H}}_{nh}$ is any particular solution of the non homogeneous Eq. (10), and $\{\vec{\mathscr{O}}_k\}$ a basis of observables spanning the null space of the associated homogeneous problem $\mathbf{X}\vec{\mathscr{H}} = 0$. The free parameters $\alpha_k \in \mathbb{R}$ can be fixed by maximizing their likelihood given sample $\mathbf{D}_{\text{eq}}$,

$$\mathscr{L}(\alpha|\mathbf{D}_{\text{eq}}) \propto \exp\left\{-\sum_{k,\mu}\alpha_k\mathscr{O}_k(\mathbf{s}^\mu) - M\log\mathscr{Z}(\alpha)\right\} \tag{13}$$

with $\mathscr{O}_k(\mathbf{s}) = (s_1,\cdots,s_N,(s_1s_2),\cdots,(s_{N-1}s_N))\cdot\vec{\mathscr{O}}_k$. The maximization provides conditions for the observables ($k = 1,...,K$),

$$\frac{\partial\mathscr{L}(\alpha)}{\partial\alpha^k} \propto \frac{1}{M}\sum_{\mu=1}^{M}\mathscr{O}_k(\mathbf{s}^\mu) - \langle\mathscr{O}_k\rangle_{\mathscr{H}} = 0\ . \tag{14}$$

Eq. (14) shows that the α_k have to be fixed such that empirical averages $\frac{1}{M}\sum_{\mu=1}^{M}\mathscr{O}_k(\mathbf{s}^\mu)$ equal model averages $\langle\mathscr{O}_k\rangle_{\mathscr{H}}$. Any possible sparsity of the matrix of measured configurations $\mathbf{X}$ (entries are 0 or 1 by definition) can be exploited to find a sparse representation of the $\{\vec{\mathscr{O}}_k\}$. In the protein example discussed above, mutagenesis experiments typically quantify all possible single-residue mutations of a reference sequence (denoted $(0,...,0)$ without loss of generality). In this case, the pairwise quantities s_is_j with $1 < i < j < N$ can be chosen as the basis $\{\mathscr{O}_k\}$ of the null space. A particular solution of the non-homogeneous system (10) is given by the paramagnetic Hamiltonian $\mathscr{H}_{nh} = \sum_i E^i s_i$, with E^i being the energy shift due to spin flip $s_i = 0 \mapsto 1$.

Artificial data

We first evaluate our method on artificial data (*Materials and Methods*). Random couplings $\mathbf{J}^0$ and fields $\mathbf{h}^0$ are chosen for a system of $N = 32$ spins. Dataset $\mathbf{D}_{\rm eq}$ is created by Markov chain Monte Carlo (MCMC) sampling, resulting in $M = 100$ equilibrium configurations. To mimic a protein 'mutagenesis' experiment, one of these configurations is chosen at random as the reference sequence, and the energies of all N configurations differing by a single spin flip from the reference (thereafter referred to as single mutants) are calculated, resulting in dataset $\mathbf{D}_E$ (after adding noise of standard deviation Δ_0). Datasets $\mathbf{D}_E$ and $\mathbf{D}_{\rm eq}$ will subsequently called "local" and "global" data respectively.

Equations (9) are solved using steepest ascent, updating parameters $\mathbf{J}$ and $\mathbf{h}$ in direction of the gradient of the joint log-likelihood (8). Since the noise Δ_0 may not be known in practical applications, we solve the equations for several values of $\lambda \in [0,1]$, weighing data sources differently. We expect the optimal inference to take place at a value λ that maximizes the likelihood in Eq. (8), *i.e.* $\lambda_0 = (1+\Delta_0^2)^{-1}$. For $\lambda = 0$, this procedure is equivalent to the classical Boltzmann machine[28], but for $\lambda > 0$, the term corresponding to the quantitative essay constrains energies of sequences in $\mathbf{D}_E$ to stay close to the measurements. As explained above, the case $\lambda = 1$ has to be treated separately; a similar gradient ascent method is used. Since exact calculations of gradients are computationally hard, the mean-field approximation is used (*Materials and Methods*).

To evaluate the accuracy of the inference, most of the existing literature on the inverse Ising modeling simply compares the inferred parameters with the true ones. However, a low error in the estimation of each inferred parameter does not guarantee that the inferred distribution matches the true one. On the contrary, in the case of a high susceptibility of the statistics with respect to parameter variations, or if the estimation of parameters is biased, the distributions of the inferred and true models could be very different even for small errors on individual parameters. For this reason, we introduce two novel evaluation procedures. First, to estimate the accuracy of the model on a local region of the configuration space, we test its ability to reproduce *energies* of configurations in $\mathbf{D}_E$. Then, we estimate the global similarity of true and inferred distributions using a measure from information theory.

Error correction of local data. We first test the ability of our approach to predict the true single-mutant energies, when noisy measurements are presented in $\mathbf{D}_E$, *i.e.* to correct the measurement noise using the equilibrium sample $\mathbf{D}_{\rm eq}$. For every λ, $\mathbf{J}$ and $\mathbf{h}$ are inferred and used to compute predicted energies of the N configurations in $\mathbf{D}_E$. The linear correlation between such predicted energies (measured with the inferred Hamiltonian) and the true energies (measured with the true Hamiltonian) is plotted as a function of λ in Fig. 2.

In the very low noise regime, $\Delta_0 \simeq 0$, the top curve in Figure 2 reaches its peak at $\lambda \simeq 1$, which is expected as local data is then sufficient to accurately "predict" energies from single mutants. On the contrary, in the high noise regime, the maximum is located close to $\lambda = 0$, pointing to the fact that local data is of little use in this case. Between those two extremes, an optimal integration strength can be found, yielding a better prediction of energies in $\mathbf{D}_{\rm E}$ as for any of the datasets taken individually. It is interesting to notice that even for highly noisy data, integrating the two sources of information with the right weight λ results in an improved modeling.

The insert of Fig. 2 shows the integration strength λ at which the best correlation is reached, against the corresponding theoretical value $\lambda_0 = 1/(1+\Delta_0^2)$ for different realizations. On average, optimal integration is reached close to the theoretical case of equation (8). This result highlights the possibility of using this integrative approach to *correct measurement errors* in the energies of single mutants. If a dataset such as $\mathbf{D}_{\rm eq}$ provides global information about the energy landscape, and the measurement noise Δ_0 can be estimated, an appropriate integration can then be used to infer more accurately the energies of the single mutants.

Global evaluation of the inferred Ising model. To assess the ability of our integrative procedure to provide a *globally* accurate model, we use the Kullback-Leibler divergence $D_{KL}(P^0||P)$ between the true model $P^0 \propto e^{-\mathscr{H}^0}$ and the inferred $P \propto e^{-\mathscr{H}}$ (*Materials and Methods*). The symmetric expression

$$\begin{aligned} \Sigma(P^0,P) &= D_{KL}(P^0||P) + D_{KL}(P||P^0) \\ &= \langle \mathscr{H}^0 - \mathscr{H} \rangle_P + \langle \mathscr{H} - \mathscr{H}^0 \rangle_{P^0} \end{aligned} \tag{15}$$

simplifies to the average difference between true and inferred energies. It can be consistently estimated using MCMC samples $\mathbf{D}$ (resp. $\mathbf{D}^0$) drawn from P (resp. P^0), without the need to calculate the partition functions. $\Sigma(P^0,P)$ has an intuitive interpretation in terms of distinguishability of models: It represents the log-odds ratio between the probability to observe samples $\mathbf{D}$ and $\mathbf{D}^0$ in their respective generating models, and the corresponding probability with models $\mathscr{H}$ and $\mathscr{H}^0$ swapped:

$$\Sigma(P^0,P) = \frac{1}{M_K} \log \left[\frac{P(\mathbf{D}|\mathscr{H}) P(\mathbf{D}^0|\mathscr{H}^0)}{P(\mathbf{D}|\mathscr{H}^0) P(\mathbf{D}^0|\mathscr{H})} \right] . \tag{16}$$

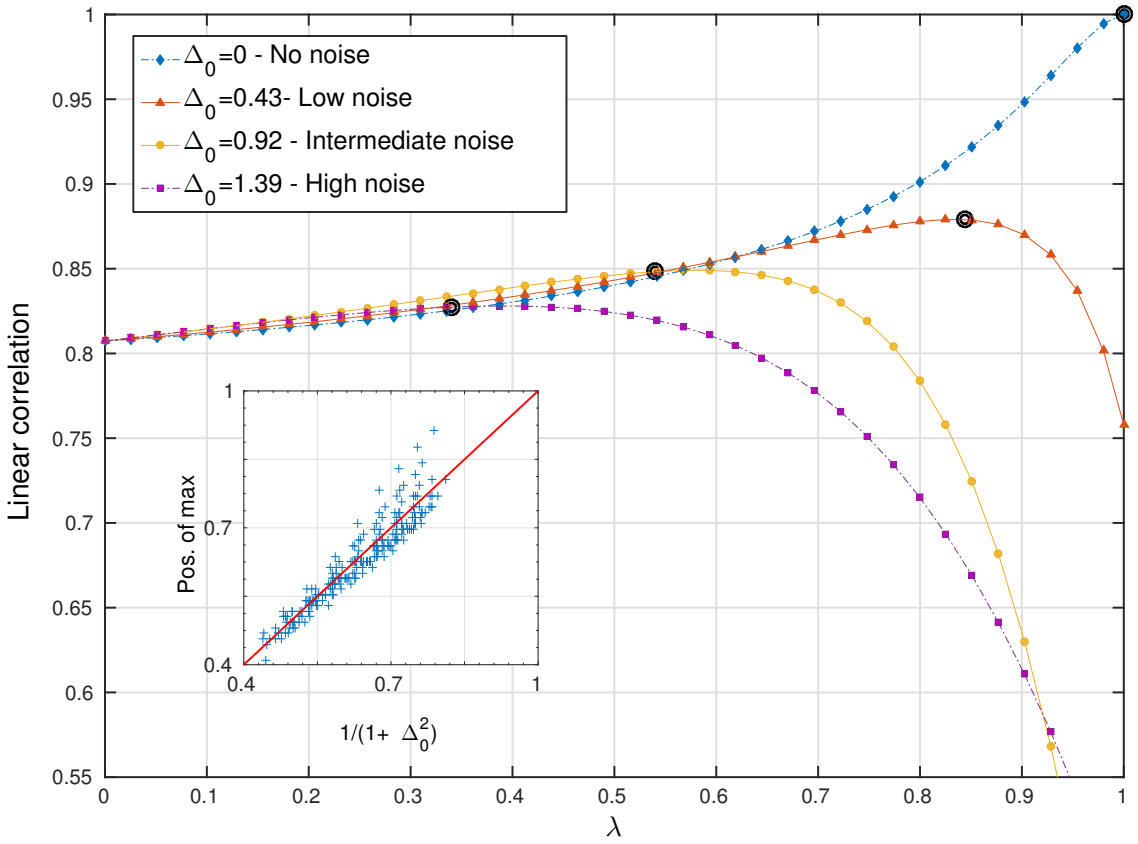

Figure 2. Integration of noisy measurements of energies of single mutants with M equilibrium configurations. The linear correlation between predicted and true single mutants energies is shown in dependence of integration parameter λ, for four different values of the noise $\Delta_0 = 0, 0.43, 0.92, 1.39$ added to the energies in the second data set. The integration strength $\lambda_0 = 1/(1+\Delta_0^2)$, which would be naturally used in case of an *a priori* known noise level, is located close to the optimal inference, cf. the black circles for $\lambda_0 = 1, 0.84, 0.54, 0.34$. The **insert** shows the value of the integration strength λ reaching maximal correlation, as a function of the theoretical value $1/(1+\Delta_0^2)$, for 200 independent realisations of the input data at different noise levels. Points are found to be closely distributed around the diagonal (red line).

where M_K is the number of sampled configurations in $\mathbf{D}^{\mathscr{H}}$ and $\mathbf{D}^0$.
The inferred model undoubtedly benefits from the integration, as a minimal divergence between the generating and the inferred probability distributions is found for an intermediate value of λ, outperforming both datasets taken individually (Figure 3). It has to be noted that even in the noiseless case $\Delta_0 = 0$, the minimum in $\Sigma(P^0, P)$ obtained at $\lambda = 1$ depends crucially on the availability of the equilibrium sample $D_{\rm eq}$. The local data D_E are not sufficient to fix uniquely all model parameters, and the degeneracy in parametrization is resolved using $D_{\rm eq}$ as explained at the end of Sec. .

As a comparison, the same analysis is done using an independent modeling that uses only fields h, and no couplings. The inset of Figure 3 clearly shows that the pairwise modeling outperforms the independent one. Even the limit $\lambda \to 1$, where $D_{\rm eq}$ becomes irrelevant in the independent model, the performance of the integrative pairwise scheme is not attained.

Biological data

To demonstrate the practical utility of our integrative framework, we apply it to the challenging problem of predicting the effect of amino-acid mutations in proteins. To do so, we use three types of data: (i) Multiple-sequence alignments (MSA) of homologous proteins containing large collections of sequences with shared evolutionary ancestry and conserved structure and function; they are obtained using HMMer[29] using profile models from the Pfam database[30]. Due to their considerable sequence divergence (typical Hamming distance $\sim 0.8N$), they provide a *global* sampling of the underlying fitness landscape. (ii) Computational predictions of the impact of all single amino-acid mutations on a protein's structural stability[31] are used to *locally* characterize the fitness landscape around a given protein. The noise term ξ^a represents the limited accuracy of this predictor, and the uncertainty in using structural stability as a proxy of protein fitness. (iii) Mutagenesis experiments have

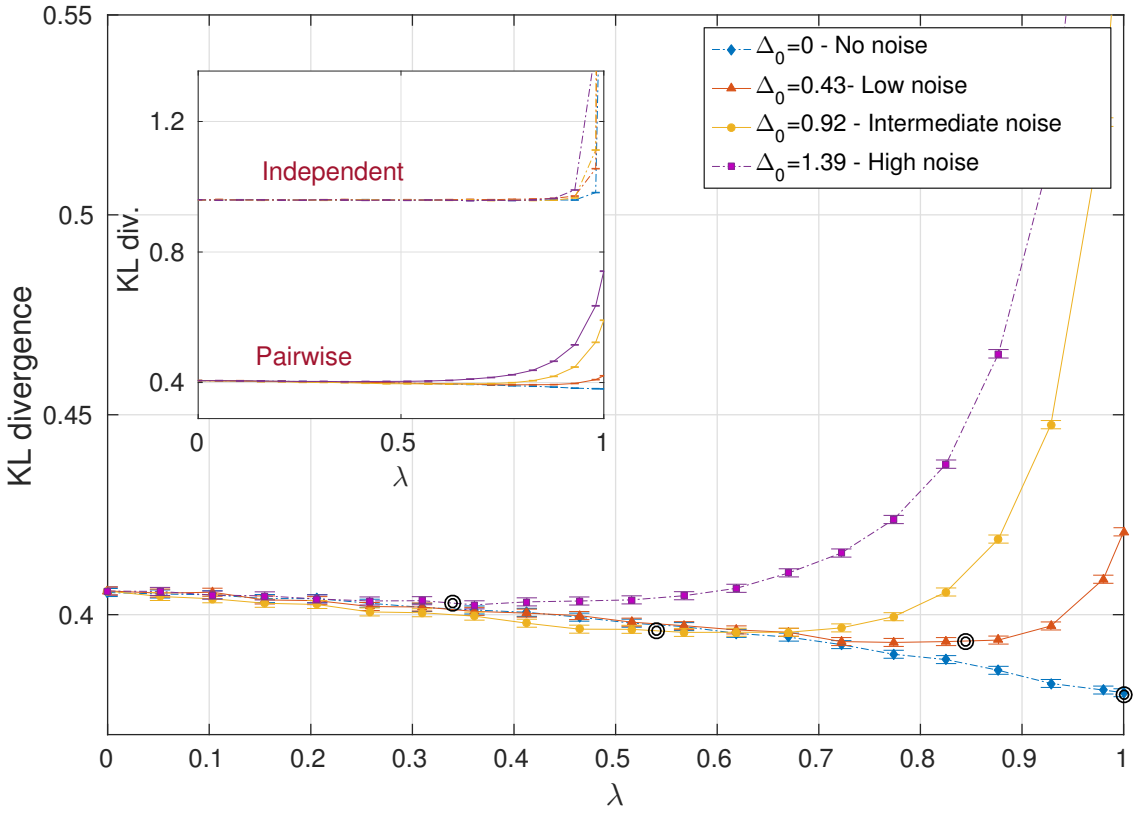

Figure 3. Symmetric Kullback-Leibler divergence Σ between true and integrated pairwise models versus strength of integration λ, estimated from $M_K = 3 \cdot 10^6$ MCMC samples. Different curves correspond to different noise levels added to single mutants energies used for integration, so that datasets $\mathbf{D}_E$ are the same as in Fig. 2. **Insert -** Comparison with an independent model using only fields $\mathbf{h}$, with the same methodology.

been used before to simultaneously quantify the fitness effects of thousands of mutants[22,23]. While datasets (i) and (ii) play the role of $\mathbf{D}_{eq}$ and $\mathbf{D}_E$ in inference, dataset (iii) is used to assess the quality of our predictions (ideally one would use the most informative datasets (i) and (iii) to have maximally accurate predictions, but no complementary dataset to test predictions would be available in that case).

To apply the inference scheme to such protein data, three modifications with respect to simulated data are needed. First, the relevant description in this case is a 21-state Potts Model (*Supporting Information*), since each variable $s_i, i = 1,...,N$, can now assume 21 states (20 amino acids, one alignment gap)[32]. Second, since measured fitnesses and model energies are found in a monotonous non-linear relation, we have used the robust mapping introduced in[18] (reviewed in the *Supporting Information*). Third, since correlations observed in MSA are typically too strong for the MF approximation to accurately estimate marginals, we relied on Markov Chain Monte Carlo (*Materials and Methods*), which has recently been shown to outperform other methods in accuracy of inference for protein-sequence data[33,34].

We have tested our approach for predicting the effect of single amino-acid mutations in two different proteins: the β-lactamase TEM-1, a bacterial enzyme providing antibiotic resistance, and the PSD-95 signaling domain belonging to the PDZ family. In both systems computational predictions can be tested against recent high-throughput experiments quantifying the *in-vivo* functionality of thousands of protein variants[22,23]. Fig. 4 shows the Pearson correlations between inferred energies and measured fitnesses as a function of the weight λ: Maximal accuracy is achieved at finite values of λ when both sources of information are combined, significantly increasing the predictive power of the models inferred considering the statistics of homologs only ($\lambda = 0$). When repeating the integrated modeling with a paramagnetic model where all sites are treated independently, $\mathscr{H}^{\text{ind}}(\mathbf{a}) = -\sum_{i=1}^{N} h_i(a_i)$ (only single-site frequencies are fitted in this case) the predictive power drops as compared to the Potts model, cf. the red lines in Fig. 4.

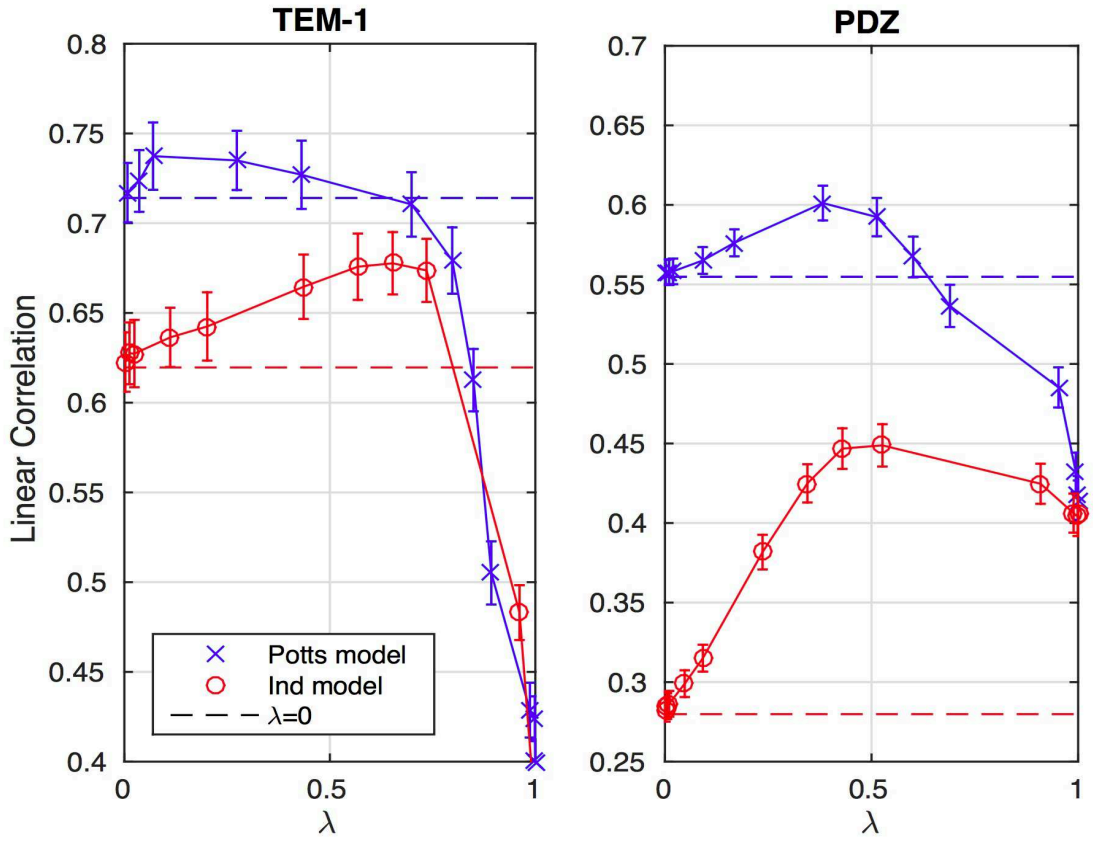

Figure 4. Linear correlation between experimental fitness effects and predictions from integrated models, at different values of λ, for 742 single mutations in the beta-lactamase TEM-1 (left panel) and for 1426 single mutations in the PSD-95 PDZ domain (right panel). Error bars represent statistical errors estimated via jack-knife estimation.

Conclusion

In this paper, we have introduced an integrative Bayesian framework for the inverse Ising problem. In difference to the standard setting, which uses only a global sample of independent equilibrium configurations to reconstruct the Hamiltonian of an Ising model, we also consider a local quantification of the energy function around a reference configuration. Using simulated data, we show that the integrated approach outperforms inference based on each single dataset alone. The gain over the standard setting of the inverse Ising problem is particularly large when the equilibrium sample is too small to allow for accurate inference.

This undersampled situation is particularly important in the context of biological data. The prediction of mutational effects in proteins is of enormous importance in various bio-medical applications, as it could help understanding complex and multi-factorial genetic diseases, the onset and the proliferation of cancer, or the evolution of antibiotic drug resistance. However, the sequence samples provided by genomic databases, like the multiple-sequence alignments of homologous proteins considered here, are typically of limited size, including even in the most favorable situations rarely more than $10^3 - 10^5$ alignable sequences. Fortunately, such sequence data are increasingly complemented by quantitative mutagenesis experiments, which use experimental high-throughput approaches to quantify the effect of thousands of mutants. While it might be tempting to use these data directly to measure mutational landscapes from experiments, it has to be noted that current experimental techniques miss at least 2-3 orders of magnitude in the number of measurable mutants to actually reconstruct the mutational landscape.

In such situations, where no single dataset is sufficient for accurate inference, integrative methods like the one proposed here will be of major benefit.

Methods

Data

Artificial data

For a system of $N = 32$ binary spins, couplings $\mathbf{J}^0$ and fields $\mathbf{h}^0$ are chosen from a Gaussian distribution with zero mean, and standard deviation $\sim 0.8/\sqrt{N}$ for J and 0.2 for h (analogous results are obtained for other parameter choices, as long as these correspond to a paramagnetic phase). Dataset $\mathbf{D}_{\rm eq}$ is created by Markov chain Monte Carlo (MCMC) simulation, resulting in $M = 100$ equilibrium configurations. A large number ($\sim 10^5$) of MCMC steps are done between each of those configurations to ensure that they are independent. One of these configurations is chosen at random as the reference sequence ("wild-type"), and the energies of all N configurations differing by a single spin flip from the reference are computed ("single mutants"). Gaussian noise of variance Δ_0^2 can be added to these energies, resulting in dataset $\mathbf{D}_E$.

Biological data

Detailed information about the analysis of biological data is provided in the *Supporting Information*.

Details of the inference

For artificial data, Eqs. (9) are solved using steepest ascent, updating parameters $\mathbf{J}$ and $\mathbf{h}$ in direction of the gradient. To ensure convergence, we have added an additional ℓ_2-regularization to the joint likelihood: $\gamma_h(\|\mathbf{h}\|_2)^2 + \gamma_J(\|\mathbf{J}\|_2)^2$. A gradient ascent method has been analogously used for the case $\lambda = 1$. To estimate the gradient, it is necessary to compute single and pair-wise probabilities $p_i(\mathbf{J},\mathbf{h})$ and $p_{ij}(\mathbf{J},\mathbf{h})$. Their exact calculation requires summation over all possible configurations of N spins, which is intractable even for systems of moderate size N, so we relied on the following approximation schemes.

Mean-field inference

In the analysis of artificial data we relied on the mean-field approximation (MF) leading to closed equations for p_i and p_{ij} :

$$\begin{aligned} p_i &= e^{h_i+\sum_{j\neq i} J_{ij} p_j}/(1+e^{h_i+\sum_{j\neq i} J_{ij} p_j}) \\ p_{ij} - p_i p_j &= -(\mathbf{J}^{-1})_{ij} \ . \end{aligned} \tag{17}$$

The main advantage of the MF approximation is its computational efficiency: The first term is solved by an iterative procedure, the second requires the inversion of the couplings matrix $\mathbf{J}$. However, the approximation is only valid and accurate at "high temperatures", *i.e.* small couplings[35]. This condition is verified in the case of the artificial data described above.

MCMC inference

Correlations observed in MSA of protein sequences are typically too strong for the MF approximation to accurately estimate marginals of the model. Therefore we use MCMC sampling of $M^{MC} = 10^4$ independent equilibrium configurations to estimate marginals at each iteration of the previously described learning protocol.

Global evaluation of the inferred Ising model

The Kullback-Leibler divergence $D_{KL}(P||Q) = \sum_s P(s)\log\{P(s)/Q(s)\}$ is a measure of the difference between probability distributions P and Q. It is zero for $P \equiv Q$, and otherwise positive. In the case of Boltzmann distributions $P \propto e^{-\mathscr{H}_P}$ and $Q \propto e^{-\mathscr{H}_Q}$, its expression simplifies to

$$D_{KL}(P||Q) = \langle \mathscr{H}_Q - \mathscr{H}_P \rangle_P + \log \mathscr{Z}_Q - \log \mathscr{Z}_P. \tag{18}$$

Evaluating this expression requires the exponential computation of the partition function of both models $\mathscr{H}_P$ and $\mathscr{H}_Q$. To overcome this difficulty, we use the symmetrized expression in Eq. (15), which only involves the average of macroscopic observables.

The symmetrized Kullback-Leibler divergence is computed by obtaining $M_K = 128\,000$ equilibrium configurations from both P and Q, using them to estimate the averages in Eq. (15).

References

1. Weigt, M., White, R. A., Szurmant, H., Hoch, J. A. & Hwa, T. Identification of direct residue contacts in protein–protein interaction by message passing. *Proceedings of the National Academy of Sciences* **106**, 67–72 (2009).
2. Mora, T., Walczak, A. M., Bialek, W. & Callan, C. G. Maximum entropy models for antibody diversity. *Proceedings of the National Academy of Sciences* **107**, 5405–5410 (2010).
3. Ferguson, A. L. *et al.* Translating hiv sequences into quantitative fitness landscapes predicts viral vulnerabilities for rational immunogen design. *Immunity* **38**, 606–617 (2013).

4. Lezon, T. R., Banavar, J. R., Cieplak, M., Maritan, A. & Fedoroff, N. V. Using the principle of entropy maximization to infer genetic interaction networks from gene expression patterns. *Proceedings of the National Academy of Sciences* **103**, 19033–19038 (2006).

5. Schneidman, E., Berry, M. J., Segev, R. & Bialek, W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* **440**, 1007–1012 (2006).

6. Cocco, S., Leibler, S. & Monasson, R. Neuronal couplings between retinal ganglion cells inferred by efficient inverse statistical physics methods. *Proceedings of the National Academy of Sciences* **106**, 14058–14062 (2009).

7. Bialek, W. *et al.* Statistical mechanics for natural flocks of birds. *Proceedings of the National Academy of Sciences* **109**, 4786–4791 (2012).

8. Jaynes, E. T. Information theory and statistical mechanics. *Physical Review* **106**, 620 (1957).

9. Roudi, Y., Tyrcha, J. & Hertz, J. Ising model for neural data: model quality and approximate methods for extracting functional connectivity. *Physical Review E* **79**, 051915 (2009).

10. Sessak, V. & Monasson, R. Small-correlation expansions for the inverse ising problem. *Journal of Physics A: Mathematical and Theoretical* **42**, 055001 (2009).

11. Mézard, M. & Mora, T. Constraint satisfaction problems and neural networks: A statistical physics perspective. *Journal of Physiology-Paris* **103**, 107–113 (2009).

12. Cocco, S., Monasson, R. & Sessak, V. High-dimensional inference with the generalized hopfield model: Principal component analysis and corrections. *Physical Review E* **83**, 051123 (2011).

13. Cocco, S. & Monasson, R. Adaptive cluster expansion for inferring boltzmann machines with noisy data. *Physical Review Letters* **106**, 090601 (2011).

14. Nguyen, H. C. & Berg, J. Mean-field theory for the inverse ising problem at low temperatures. *Physical Review Letters* **109**, 050602 (2012).

15. Aurell, E. & Ekeberg, M. Inverse ising inference using all the data. *Physical Review Letters* **108**, 090201 (2012).

16. Nguyen, H. C. & Berg, J. Bethe–peierls approximation and the inverse ising problem. *Journal of Statistical Mechanics: Theory and Experiment* **2012**, P03004 (2012).

17. Decelle, A. & Ricci-Tersenghi, F. Pseudolikelihood decimation algorithm improving the inference of the interaction network in a general class of ising models. *Physical Review Letters* **112**, 070603 (2014).

18. Figliuzzi, M., Jacquier, H., Schug, A., Tenaillon, O. & Weigt, M. Coevolutionary inference of mutational landscape and the context dependence of mutations in beta-lactamase tem-1. *Molecular Biology and Evolution* (2016).

19. Asti, L., Uguzzoni, G., Marcatili, P. & Pagnani, A. Maximum-entropy models of sequenced immune repertoires predict antigen-antibody affinity. *PLoS Comput Biol* **12**, e1004870 (2016).

20. Morcos, F., Schafer, N. P., Cheng, R. R., Onuchic, J. N. & Wolynes, P. G. Coevolutionary information, protein folding landscapes, and the thermodynamics of natural selection. *Proceedings of the National Academy of Sciences* **111**, 12408–12413 (2014).

21. Mann, J. K. *et al.* The fitness landscape of hiv-1 gag: Advanced modeling approaches and validation of model predictions by in vitro testing. *PLoS Comput Biol* **10**, e1003776 (2014).

22. McLaughlin Jr, R. N., Poelwijk, F. J., Raman, A., Gosal, W. S. & Ranganathan, R. The spatial architecture of protein function and adaptation. *Nature* **491**, 138–142 (2012).

23. Jacquier, H. *et al.* Capturing the mutational landscape of the beta-lactamase tem-1. *Proceedings of the National Academy of Sciences* **110**, 13067–13072 (2013).

24. Melamed, D., Young, D. L., Gamble, C. E., Miller, C. R. & Fields, S. Deep mutational scanning of an rrm domain of the saccharomyces cerevisiae poly (a)-binding protein. *RNA* **19**, 1537–1551 (2013).

25. Hinkley, T. *et al.* A systems analysis of mutational effects in hiv-1 protease and reverse transcriptase. *Nature genetics* **43**, 487–489 (2011).

26. de Visser, J. A. G. & Krug, J. Empirical fitness landscapes and the predictability of evolution. *Nature Reviews Genetics* **15**, 480–490 (2014).

27. Otwinowski, J. & Plotkin, J. B. Inferring fitness landscapes by regression produces biased estimates of epistasis. *Proceedings of the National Academy of Sciences* **111**, E2301–E2309 (2014).

28. Ackley, D. H., Hinton, G. E. & Sejnowski, T. J. A learning algorithm for boltzmann machines. *Cognitive Science* **9**, 147–169 (1985).

29. Mistry, J., Finn, R. D., Eddy, S. R., Bateman, A. & Punta, M. Challenges in homology search: Hmmer3 and convergent evolution of coiled-coil regions. *Nucleic Acids Research* **41**, e121 (2013).

30. Finn, R. D. *et al.* Pfam: the protein families database. *Nucleic Acids Research* **42**, D222–D230 (2014).

31. Dehouck, Y., Kwasigroch, J. M., Gilis, D. & Rooman, M. Popmusic 2.1: a web server for the estimation of protein stability changes upon mutation and sequence optimality. *BMC Bioinformatics* **12**, 151 (2011).

32. Morcos, F. *et al.* Direct-coupling analysis of residue coevolution captures native contacts across many protein families. *Proceedings of the National Academy of Sciences* **108**, E1293–E1301 (2011).

33. Sutto, L., Marsili, S., Valencia, A. & Gervasio, F. L., From residue coevolution to protein conformational ensembles and functional dynamics. *Proceedings of the National Academy of Sciences* **112**, 13567–13572 (2015).

34. Haldane, A., Flynn, W. F., He, P., Vijayan, R. S. K. & Levy, R. M., Structural Propensities of Kinase Family Proteins from a Potts Model of Residue Co-Variation. *Protein Science* **25**, 1378–1384 (2016).

35. Plefka, T. Convergence condition of the tap equation for the infinite-ranged ising spin glass model. *Journal of Physics A: Mathematical and general* **15**, 1971 (1982).

Acknowledgements

MW acknowledges funding by the ANR project COEVSTAT (ANR-13-BS04- 0012-01). This work undertaken partially in the framework of CALSIMLAB, supported by the grant ANR-11-LABX-0037-01 as part of the "Investissements d'Avenir" program (ANR-11-IDEX-0004-02).

Author contributions statement

M.F. and M.W. designed research; P.B.C., M.F. and M.W. performed research; P.B.C. and M.F. analyzed data; and P.B.C., M.F. and M.W. wrote the paper.

Additional information

The author(s) declare no competing financial interests.